

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: REMOTELY CONTROLLING A UNIX-BASED SYSTEM
APPLICANT: ERIC D. FAGERBURG AND CHARLES M. MORGAN

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL68839084US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

12-28-00
Date of Deposit

Derek W. Norwood
Signature

Derek W. Norwood
Typed or Printed Name of Person Signing Certificate

REMOTELY CONTROLLING A UNIX-BASED SYSTEM

BACKGROUND

This invention relates to remotely controlling a UNIX-based system.

UNIX is a multi-user, multitasked operating system typically written in the C programming language. Any machine capable of running (compiling) programs written in C (or whatever programming language is used) can install and run UNIX. Users can enter UNIX commands on such a machine for execution.

DESCRIPTION OF DRAWINGS

FIG. 1 is a diagram of a network configuration.

FIG. 2 is a flowchart showing a process of remotely controlling a UNIX-based system.

FIG. 3 shows a terminal displaying a current screen.

FIG. 4 shows the terminal displaying the current screen of FIG. 3 plus a prompt.

FIG. 5 shows a terminal displaying a prompt.

DESCRIPTION

Referring to FIG. 1, a simplified network configuration includes a user 102 at a user terminal 104 running a UNIX-based system. The user 102 has the ability to grant or

deny permission to a remote user 106 at a remote terminal 108 to remotely control the user terminal 104 from the remote terminal 108.

When the remote user 106 wants to remotely control the user terminal 104, a remote viewer agent 112 at the remote terminal 108 notifies a remote control agent 110 at the user terminal 104. The remote control agent 110 runs as a UNIX-based daemon running in the background of any programs or applications currently running on the user terminal 104. The remote control agent 110 can be included as part of the user terminal 104, e.g., as a stored program, or be otherwise accessible to the user terminal 104, e.g., across the network 118.

Similarly, the remote viewer agent 112 can be included as part of the remote terminal 108 or be otherwise accessible to the remote terminal 108. The remote viewer agent 112 can be used to control any number of user terminals across any network, not just the shown network 118. For example, the remote terminal 108 may include or have access to a collection of data 124 including a list of user terminals connected to a network. The remote user 106 can choose to remotely control any of those user terminals over the network.

The remote viewer agent 112 sends notification to the remote control agent 110 over a network 118 and communication

links 120a and 120b. The communication links 120a and 120b between the user terminal 104 and the remote terminal 108 can be any kind and any combination of communication links such as modem links, cables, point-to-point links, infrared
5 connections, fiber optic links, cellular links, Bluetooth, satellite links, or other similar types of links.

Furthermore, the network 118 can include any kind and any combination of local networks, private networks, public networks, or other similar types of networks.

10 The remote agent 110 creates a new session at the user terminal 104 and allocates a new virtual terminal 126 in the background of the user terminal 104. The remote agent 110 duplicates the current contents displayed on a display screen 122 of the user terminal 104 into the new virtual terminal
15 126, inserts a prompt on the new virtual terminal 126, and switches the display on the display screen 122 to the new virtual terminal display 126 (the current contents plus the prompt).

When a machine is running in text mode, a daemon
20 typically does not interact with the machine's console (display terminal or screen), i.e., the daemon typically closes stdin, stdout, and stderr. No messages are shown on the console indicating when or if the daemon has taken any action. Instead, the daemon's actions are recorded in a

collection of data for the machine, such as a log file, that can be checked after-the-fact to determine the daemon's activities (if any). By displaying the text prompt to the user 102, the remote control agent 110 lets the user 102 know
5 in real time when the remote user 108 wants to take remote control of the user terminal 104. Furthermore, the user 102 has the ability to allow or prevent the remote control of the user terminal 104 by the remote user 106.

Generally, the prompt asks the user if a remote user
10 should be allowed to control the user terminal 104. The user 102 can enter a response to the prompt using, for example, a keyboard 114 or a mouse 116. If the user 102 grants permission, then the remote user 106 can control the user terminal 104 as if the remote user 106 was directly using the
15 user terminal 104. For example, the remote user 106 can issue text commands to the user terminal 104 via the network 118 and the communication links 120a and 120b. Actions that the remote user 106 may issue to the user terminal 104 include transferring files, deleting files, adding files, rebooting
20 the user terminal 104, and other similar actions allowable in the UNIX-based system.

The user terminal 104 can be any machine capable of running (compiling) programs written in the same programming language (e.g., C, machine language, etc.) as the UNIX-based

system and that can install and run the UNIX-based system, such as portable and stationary computers, personal digital assistants, network servers, and other similar devices. Similarly, the user terminal 104 can use any hardware platform, such as disk operating system (DOS), a Microsoft Windows-based system, a Macintosh-based system, a Linux-based system, or other similar platform.

The user terminal 104 can use any version of UNIX, such as Linux, System III (UNIX by AT&T, American Telephone & Telegraph Company), System V (UNIX by AT&T), Solaris (UNIX by Sun Microsystems, Inc.), HP-UX (UNIX by Hewlett-Packard Company), Advanced Interactive Executive (AIX, UNIX by IBM, International Business Machines, Inc.), UnixWare (UNIX by SCO, the Santa Cruz Operation, Inc.), Berkeley UNIX versions, and other similar UNIX versions.

Referring to FIG. 2, the remote user 106 can attempt to remotely control the user terminal 104 using, for example, a remote control process 200. The remote control process 200 starts 202 when the remote viewer agent 112 notifies 204 the remote control agent 110 that the remote terminal 108 wants to remotely control the user terminal 104. The remote viewer agent 112 sends this notification when the remote user 106 enters a request at the remote terminal 108 to remotely

control the user terminal 104 or otherwise engages the remote viewer agent 112.

The remote control agent 110 determines 206 if the user terminal 104 is running in text mode. Text mode is a screen display mode that only displays text; it does not display graphics. In other words, if a machine is operating in text mode, it only displays characters, not pictures or graphics images.

If the user terminal 104 is not running in text mode, then the remote control agent 110 presumes that the user terminal 104 can display graphics and presents 208 a graphics-based permission prompt on the user terminal 104 to the user 102. For example, the prompt may be a dialog box with a "give permission" button and a "refuse permission" button that the user 102 can click on and select with the mouse 116. Based on the user's response to the prompt, the remote user 106 may be able to remotely control the user terminal 104.

If the user terminal 104 is running in text mode, then the remote control agent 110 creates 212 a new session and the virtual console 126 in the background of the user terminal's currently active console. In other words, the remote control agent 110 creates or compiles the instructions necessary to create a new screen for display on the user terminal 104 on

the display screen 122. This new screen is created or compiled transparently to the user 102, i.e., without the user's knowledge. For example, the remote control agent 110 may open a new tty (here, a new, "virtual" terminal) in
5 non-blocking mode and create a new stdin (a file pointer that provides access to the keyboard 114) and a new stdout and a new stderr (file pointers that write data to the display screen 122) for the new tty.

The remote control agent 110 replicates 214 the contents
10 of the user terminal's currently active display screen 122 into the virtual console 126. To replicate the current contents of the user terminal's display screen 122, the remote control agent 110 gathers data about the display screen 122 such as the size of the display screen 122 and the cursor's
15 position on the display screen 122. The remote control agent 110 also determines and reserves the amount of memory necessary to store the current contents of the display screen 122 to use in the replication.

The remote control agent 110 also inserts 216 a text
20 prompt into the virtual console 126. The text prompt appears as the last text line on the virtual console 126 (although it could be anywhere) and informs the user 102 that a remote user would like to remotely control the user terminal 104. For example, the text prompt could be "Will you allow <remote

user> to remote control your machine? Y/N," "Will you allow
<remote user> to transfer files? Y/N," "May another user
remote control your machine?," or other similar message.

One or more text prompts may be available for display by
5 the remote control agent 110. If two or more text prompts are
available, the remote control agent 110 may choose which text
prompt to display to the user 102 based on information sent in
the notification from the remote terminal 108. For example,
the notification could indicate that the remote user 108
10 desires to transfer files from the user terminal 104, in which
case the remote control agent 110 would choose a text prompt
such as "May <remote user> transfer files? Y/N."

After creating the virtual console 126 and inserting the
text prompt, the remote control agent 110 switches 218 from
15 the currently active display screen 122 (shown for example in
FIG. 3) to the virtual console 126 (shown for example in FIG.
4). In this way, the user 102 sees on the display screen 122
what he or she was seeing before the switch to the virtual
console 126 (except for the addition of the text prompt). The
20 user 102 experiences minimal visual disruption during the
switching, likely not being able to detect that any virtual
terminal switching was done (except for the addition of the
prompt).

Instead of inserting the text prompt into the virtual console 126, the remote control agent 110 may copy and store the current contents of the display screen 122 but display a screen with only the text prompt as shown for example in FIG.

5 5.

The remote control agent 110 may also switch the user terminal 104 into raw mode so the remote control agent 110 can more quickly receive user input. In raw mode, data input to the user terminal 104, e.g., a response to the prompt, is passed to the appropriate device or mechanism, e.g., the remote control agent 110, before being otherwise processed or interpreted.

The remote control agent 110 determines 220 if the user 102 responded to the text prompt. If the user 102 responded, then the user's response is returned 222 to the remote control agent 110. If the user 102 has not responded, e.g., entered nothing or entered data not in a form recognizable by the remote control agent 110 such as the recognizable "Y" or "N", "y" or "n", "1" or "2", "yes" or "no", etc., then the remote control agent 110 determines 224 if a threshold amount of time has passed since the text prompt was displayed to the user 102. The threshold amount of time is a predetermined time interval in which the user must respond to the text prompt before a default response is returned to the remote control

agent 110. If the threshold amount of time has not expired, then the user 102 is given another text prompt requesting input, either giving the same text prompt again or asking in a different text prompt format for another response because the previous response was in an unknown format. If the threshold amount of time has not expired, then a default response is returned 226 to the remote control agent 110. The default response can be either to allow or to disallow remote control of the user terminal 104.

Once the remote control agent 110 receives a response to the text prompt, the remote control agent 110 switches 228 the virtual console 126 back to the user's original display screen 122 (the currently active display screen 122 that the user 102 saw before the user terminal 104 was switched to the virtual console 126) and deallocates the virtual console 126. The deallocation can also include deallocating any memory reserved to store the contents of the currently active console.

The remote control agent 110 also determines 230 if the response to the text prompt allows the remote user 106 to remotely control the user terminal 104. If permission was granted by the user 102 (or through the default response), then the remote control agent 110 informs 232 the remote user 106, who can then remotely control the user terminal 104. If permission was not granted, then the remote control agent 110

informs 234 the remote user 106, who may not now remotely control the user terminal 104 but may later request to remotely control the user terminal 104.

The techniques described here are not limited to any particular hardware or software configuration; they may find applicability in any computing or processing environment. The techniques may be implemented in hardware, software, or a combination of the two. Preferably, the techniques are implemented in programs executing on programmable machines such as mobile or stationary computers, personal digital assistants, and similar devices that each include a processor, a storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and one or more output devices. Program code is applied to data entered using the input device to perform the functions described and to generate output information. The output information is applied to one or more output devices.

Each program is preferably implemented in a high level procedural or object oriented programming language to communicate with a machine system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language.

Each such program is preferable stored on a storage medium or device, e.g., compact disc read only memory (CD-ROM), hard disk, magnetic diskette, or similar medium or device, that is readable by a general or special purpose programmable machine for configuring and operating the machine when the storage medium or device is read by the computer to perform the procedures described in this document. The system may also be considered to be implemented as a machine-readable storage medium, configured with a program, where the storage medium so configured causes a machine to operate in a specific and predefined manner.

Other embodiments are within the scope of the following claims.